ANITA proposal for a National Astronomy computing strategy

Daniel Price (ANITA Chair), Ashley Ruiter, Ivo Seitenzahl, Stas Shabala, J-P Macquart, Roland Crocker, Rhys Poulton, Jacob Seiler (ANITA steering committee), with input from Christoph Federrath, Mark Krumholz, Simon Campbell, Chris Power, Jonti Horner, Bernhard Mueller, Alexander Heger, Tamara Davis and Jarrod Hurley, on behalf of ANITA membership.

22nd October 2018

Introduction

In a letter to AAL dated 29th March 2018 the ANITA steering committee expressed concern about the implementation of the decadal plan priority for "World-class high performance computing (HPC) and software capability for large theoretical simulations, and resources to enable processing and delivery of large data sets from these facilities".

ANITA is concerned that this priority has so far been narrowly interpreted as "delivery of large data sets", i.e., as a *data strategy*, while a *computing strategy* is lacking. The Decadal Plan (p49) outlined a High Performance Computing (HPC) priority in simple terms as "access equivalent to 30% of a top-100 supercomputer".

As quoted in the letter: "ANITA supports the idea of Virtual Observatories, and AAL funding for them, but it is the informed consensus of the ANITA committee that funding for ASVO nodes should not be seen as the primary fulfilment of the above Decadal Plan priority". We expressed concern that the priority the theory community expressed in the Australian Astronomy Decadal Plan was being misinterpreted.

Requirements for a computing strategy

The key needs of the theory community are:

Flexibility: Theory is not "one size fits all". A range of resources are needed to satisfy a range of different computational workflows. We need access to resources appropriate to each workflow. These range from large individual jobs with tightly coupled networking requirements, to small simulations running on tens to hundreds of CPUs, to embarassingly parallel parameter searches with little or no communication between tasks.

High throughput computing: One of the most important metrics for the usefulness of a computing facility is time-to-begin-calculation — the ability to get compute jobs to

start within 24 hours, which is related to a) having dedicated facilities available b) machine uptime and c) user support. The OzSTAR facility has been particularly successful in this respect. Interruptions to throughput such as the recent disk failure on the ageing gSTAR or downtime for maintenance or updates are particularly fatal to productivity.

Access to an internationally competitive level of resources: Australia currently only has one machine in the top 100 supercomputers (Raijin at NCI), followed by Magnus (Pawsey) ranked #217. Only 115M cpu-hours is available on Raijin as open time through NCMAS2019 for the *entire* scientific community[1]. Only a fraction of this is available to Astronomy, 1-2 orders of magnitude less than the goal in the Decadal Plan. This compares to *individual* allocations of upwards of 100M cpu-hours available through top-level European facilities. By contrast, NCMAS considers applications of > 4M cpu-hours `large'.

Low barrier to entry: Good young researchers move between countries and universities, new students begin every year, and each require rapid access to compute facilities appropriate to their workflow. Ease of access to OzSTAR via "join the queue" access is one of the main reasons this machine is highly valued by research groups. Similar arrangements, e.g. by partner share of ANU or UNSW with the NCI facility are highly valued for the same reason. Compute workflows often cannot be planned 12 months in advance (e.g. by application via NCMAS), as they may require significant code development or student training.

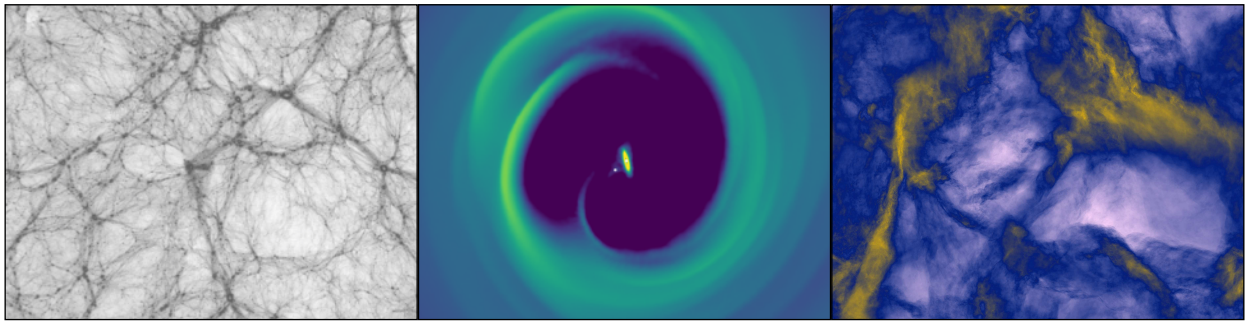

Highlights from recent theoretical and computational astrophysics research in Australia. Left: Synthetic UniveRses for Surveys (SURFS) cosmological simulations (Elahi et al. 2018). Middle: Simulations of a young circumbinary disc (Price et al. 2018). Right: Snapshot from the world's highest-resolution simulation of turbulence (10048^3 resolution elements, Federrath et al. 2016).

Structuring of resources

ANITA recommends a structured provision of computing resources to the Australian Astronomical community. We recommend a tiered categorisation of resources defined by the level of parallel computing infrastructure required, similar to the categorisation used by the European Union's PRACE strategy:

[1] https://ncmas.nci.org.au/2019/sites/default/files/public/pdfs/NCMAS-2019-InformationForApplicants.pdf

Tier 0: National level machines at the level of a Top 100 supercomputer. Jobs typically use thousands of CPUs. Access by merit allocation. A typical Astronomy user would be able to request an individual annual allocation of 10-100 million CPU hours.

Tier 1: Jobs that use tens to hundreds of CPU cores, including jobs that use shared-memory parallelisation but no distributed memory. Includes machines with specialised hardware such as Knights Landing or GPUs. A merit allocation would be typically award 1-10 million CPU hours to individual Astronomy users annually.

Tier 2: Jobs that require many single CPUs but no communication between tasks. These may use hundreds or even thousands of CPUs. Access to such machines needs flexibility, i.e., lots of jobs run in a short fraction of time, which may be best supported on cloud computing facilities. Such workflows should be not allowed to clog Tier 0 or Tier 1 facilities. A typical Astronomy user would be able to access at least 10 million CPU hours annually, equating to continuous use of around 1,000 CPUs.

Tier level	Description	Typical workflow	Interconnect	Example facilities	Astronomy applications	Access mechanism
Tier 0	National centres open to all disciplines	Large individual jobs running on thousands of CPUs	Tightly coupled / infiniband or similar	NCI, Pawsey	Cosmological simulations, supernovae or star formation simulations using MPI parallel fluids codes	Merit allocation scheme, ideally twice per year
Tier 1	National level but specialised facilities (e.g. specific to Astronomy)	Parameter studies using tens to hundreds of CPUs per job	Tightly coupled / infiniband or similar	gSTAR, OzSTAR, Astronomy partner share in national facilities	Fluid simulations using MPI / OpenMP parallel codes. N-body calculations using GPUs.	Available on demand, with large allocations by merit allocation
Tier 2	Grid / cloud computing	Large number of single CPU jobs	No communication between jobs	NeCTAR grid	Parameter estimation, Monte-Carlo Markov chain, Bayesian inference, Stellar evolution codes	Available on demand

The missing pieces

We identify two missing pieces in the current Astronomy computing provision: i) There are effectively no internationally competitive Tier-0 allocations available to astronomy and ii) We currently lack a Tier 2 facility that is easily accessible or widely used by the community. Access for Astronomers needs to be configured in such a way that it can be accessed in the same manner as a Tier-1 facility, i.e. via a job submission queue.

ANITA recommendations for a National Astronomy Computing Strategy

1. **We recommend AAL either fund, or secure the use of (via NeCTAR or similar), the continuous use of (on average) 5,000 single VCPUs as an Astronomy-dedicated Tier-2 cloud computing facility** for parameter estimation, Bayesian inference and other `embarassingly parallel' single-core computational tasks. These resources should be accessible via a job queueing system (e.g., SLURM queue) similar to how users submit jobs to other national facilities. Estimated cost based on Amazon EC2 pricing ($0.0132 per VCPU-hour for t3.small) is $50k/month, but this may also be possible to achieve more cost-effectively in partnership with NeCTAR. The advantage of a cloud service is that billing is for actual usage only. This would in turn relieve congestion on OzSTAR allowing for better use of this Tier-1 facility for compute jobs requiring shared and distributed memory parallelism (OpenMP/MPI) and/or GPUs.
2. **We recommend planning for an upgrade of the OzSTAR facility** on a 3-5 year timescale, with hardware investment comparable to AAL's initial $1m investment in gSTAR, keeping this as Australian Astronomy's primary Tier-1 facility.
3. **We recommend AAL lobby towards securing access to an internationally competitive Tier-0 facility** with 200M cpu hours of time available for Astronomy annually. This could be achieved, for example, by Australia joining the EU PRACE scheme, or by major upgrades to existing national facilities.
4. **We recommend merit allocation calls occur twice per year, or on a rolling basis.** Semi-annual calls are common elsewhere (e.g. DiRAC, Juelich, PRACE).
5. **We recommend prioritising funding towards software support staff** as part of ADACS. Specifically we would suggest that ADACS priorities should shift away from the provision of generic training courses, towards more tailored support for researchers, including help to shift workflows to appropriate tiers. We see the availability of such software support staff as an excellent way of satisfying the decadal plan priority for software capability for large theoretical simulations. These staff are also a unique strength of ADACS and the Australian astronomy community.

Key to terminology:

ANITA=Australian National Institute for Theoretical Astrophysics (anita.edu.au)

AAL=Astronomy Australia Limited (astronomyaustralia.org.au)

ASVO=All-Sky Virtual Observatory (asvo.org.au)

HPC=High Performance Computing

CPU=Central Processing Unit; VCPU=Virtual CPU

EC2=Elastic Cloud 2 (aws.amazon.com/ec2/)

NeCTAR=National Collaboration Tools And Resources project (nectar.org.au)

ADACS=Astronomy Data and Computing Service (adacs.org.au)

ANU=Australian National University

UNSW=University of New South Wales

PRACE=Partnership for Advanced Computing in Europe (prace-ri.eu)

OpenMP=Open Multi-Processing (openmp.org); MPI=Message Passing Interface (mpi-forum.org; open-mpi.org)

NCMAS=National Compute infrastructure Merit Allocation Scheme